

ITS Technical Bulletin 0283
DB2 Version 6 Release 1

Issued Date: June 5, 2000
Effective Date: June 25, 2000
Section/Groups: Software Management/Database Administration
Submitted By: Vicki Stokes
Approved By: Jim Calaway

On Sunday, June 25, 2000, the State of Utah will install DB2 Version 6 Release 1 on CPU4 in the DBD1 development subsystem. This software will impact ORSIS WERD, WERI, WERA, and WERT regions. Subsequent installation on CPU3 is scheduled for August 2000 and will impact ORSIS production WERP.

All DB2 PLANS and PACKAGES will be rebound and PLAN_TABLES will be altered for Version 6.

New version features can be exploited after the time when fallback to Version 5 would not be a consideration.

Summary of Changes to DB2 UDB Server for OS/390, Version 6
(from IBM Installation Guide)

DB2 UDB for OS/390 Version 6 delivers an enhanced relational database server solution for OS/390. This release focuses on greater capacity, performance improvements for utilities and queries, easier database management, more powerful network computing, and DB2 family compatibility with rich new object-oriented capability, triggers, and more built-in functions.

Capacity Improvements

16-terabyte tables provide a significant increase to table capacity for partitioned and LOB table spaces and indexes, and for nonpartitioning indexes.

Buffer pools in data spaces provide virtual storage constraint relief for the ssnmDBM1 address space, and data spaces increase the maximum amount of virtual buffer pool space allowed.

Performance and Availability

Improved partition rebalancing lets you redistribute partitioned data with minimal impact to data availability. One REORG of a range of partitions both reorganizes and rebalances the partitions.

You can change checkpoint frequency dynamically using the new SET LOG command and initiate checkpoints any time while your subsystem remains available.

Utilities that are faster, more parallel, easier to use:

Faster backup and recovery enable COPY and RECOVER to process a list of objects in parallel, and recover indexes and table spaces at the same time from image copies and the log.

Parallel index build reduces the elapsed time of LOAD and REORG jobs of table spaces, or partitions of table spaces, that have more than one index; the elapsed time of REBUILD INDEX jobs is also reduced.

Tests show decreased elapsed and processor time for online REORG.

Inline statistics embeds statistics collection into utility jobs, making table spaces available sooner.

You can determine when to run REORG by specifying threshold limits for relevant statistics from the DB2 catalog.

Query performance enhancements include:

Query parallelism extensions for complex queries, such as outer joins and queries that use nonpartitioned tables

Improved workload balancing in a Parallel Sysplex® that reduces elapsed time for a single query that is split across active DB2 members.

Improved data transfer that lets you request multiple DRDA query blocks when performing high-volume operations.

The ability to use an index to access predicates with noncorrelated IN subqueries

Faster query processing of queries that include join operations.

More performance and availability enhancements include:

Faster restart and recovery with the ability to postpone backout work during restart, and a faster log apply process.

Increased flexibility with 8-KB and 16-KB page sizes for balancing different workload requirements more efficiently, and for controlling traffic to the coupling facility for some

workloads.

Direct-row access using the new ROWID data type to re-access a row directly without using the index or scanning the table.

Ability to retain prior access path when you rebind a statement. You almost always get the same or a better access path. For the exceptional cases, Version 6 of DB2 for OS/390 lets you retain the access path from a prior BIND by using rows in an Explain table as input to optimization.

An increased log output buffer size (from 1000 4-KB to 100000 4-KB buffers) that improves log read and write performance.

Data Sharing Enhancements

More caching options use the coupling facility to improve performance in a data sharing environment for some applications by writing changed pages directly to DASD.

Control of space map copy maintenance with a new option avoids tracking of page changes, thereby optimizing performance of data sharing applications.

User Productivity

Predictive governing capabilities enhance the resource limit facility to help evaluate resource consumption for queries that run against large volumes of data.

Statement cost estimation of processing resource that is needed for an SQL statement helps you to determine error and warning thresholds for governing, and to decide which statements need tuning.

A default buffer pool for user data and indexes isolates user data from the DB2 catalog and directory, and separating user data from system data helps you make better tuning decisions.

More information available for monitoring DB2 includes data set I/O activity in traces, both for batch reporting and online monitors.

Better integration of DB2 and Workload Manager delay reporting enables DB2 to notify Workload Manager about the current state of a work request.

More tables are allowed in SQL statements SELECT, UPDATE, INSERT, and DELETE, and in views. DB2 increases the limit from 15 to 225 tables. The number of tables and views in a subselect is not changed.

Improved DB2 UDB family compatibility includes SQL extensions, such as:

- A VALUES clause of INSERT that supports any expression
- A new VALUES INTO statement

Easier recovery management lets you achieve a single point of recovery and recover data at a remote site more easily.

Enhanced database commands extend support for pattern-matching characters (*) and let you filter display output.

You can easily process dynamic SQL in batch mode with the new object form of DSNTEP2 shipped with DB2 for OS/390.

Network Computing

SQLJ, the newest Java implementation for the OS/390 environment, supports SQL embedded in the Java programming language. With SQLJ, your Java programs benefit from the superior performance, manageability, and authorization available to static SQL, and they are easy to write.

DRDA® support for three-part names offers more functionality to applications using three-part names for remote access and improves the performance of client/server applications.

Stored procedure enhancements include the ability to create and modify stored procedure definitions, make nested calls for stored procedures and user-defined functions, and embed CALL statements in application programs or dynamically invoke CALL statements from IBM's ODBC and CLI drivers.

DB2 ODBC extensions include new and modified APIs and new data types to support the object-relational extensions.

ODBC access to DB2 for OS/390 catalog data improves the performance of your ODBC catalog queries by redirecting them to shadow copies of DB2 catalog tables.

Better performance for ODBC applications reduces the number of network messages that are exchanged when an application executes dynamic SQL.

Improvements for dynamically prepared SQL statements include a new special register that you use to implicitly qualify names of distinct types, user-defined functions, and stored procedures.

DDF connection pooling uses a new type of inactive thread that improves performance for large volumes of inbound DDF connections.

Object-relational Extensions and Active Data

Large objects (LOBs) are well suited to represent large, complex structures in DB2 tables. Now you can make effective use of multimedia by storing objects such as complex documents, videos, images, and voice.

Distinct types (which are sometimes called user-defined data types), like built-in data types, describe the data that is stored in columns of tables where the instances (or objects) of these data types are stored. They ensure that only those functions and operators that are explicitly defined on a distinct type can be applied to its instances.

User-defined functions, like built-in functions or operators, support manipulation of distinct type instances (and built-in data types) in SQL queries.

New and extended built-in functions improve the power of the SQL language with about 100 new built-in functions, extensions to existing functions, and sample user-defined functions.

Triggers automatically execute a set of SQL statements whenever a specified event occurs. These statements validate and edit database changes, read and modify the database, and invoke functions that perform operations inside and outside the database.

You can use the DB2 Extenders feature of DB2 for OS/390 to store and manipulate image, audio, video, and text objects. The extenders automatically capture and maintain object information and provide a rich body of APIs. **Note:** Extenders not currently installed see ITS DBAs if you need this feature.

Features of DB2 for OS/390

DB2 for OS/390 Version 6 offers a number of tools, which are optional features of the server, that are shipped to you automatically when you order DB2 Universal Database for OS/390:

- DB2 Management Tools Package, which includes the following elements:
 - DB2 UDB Control Center
 - DB2 Stored Procedures Builder - available soon
 - DB2 Installer
 - DB2 Visual Explain
 - DB2 Estimator
 - Net.Data® for OS/390 - currently on cpu0

- Query Management Facility 6.1
- DB2 DataPropagator(TM) Capture 6.1

QMF 6.1 and Data Propagator Capture 6.1 will be migrated approximately two to three weeks after DB2 V6 has been installed. The 6.1 release of these products exploit version 6 features of DB2.

C and C++ compile procedures are available in SYSP.PROCLIB.

DBD1C	-	DBP1C	C	Compile Procedure
DBD1CPP	-	DBP1CPP	CPP	Compile Procedure